

Game Playing in AI Research and Education system

Pranav Kumar Pathak, Dr. Dipthi Shah, PhD (CS)

ABSTRACT: *Introduced in 2005 as a new AI Challenge and Competition, general game playing has quickly evolved into an established research area. More recently it is also gaining popularity as a useful addition to AI curricula at universities around the world. The first part of this paper will survey the research landscape of general game playing, which covers a broad range of classic AI topics, including knowledge representation, search, planning and learning. The second part will argue that general game playing provides a unique approach to teaching a number of different topics such as problem solving by search, logic, logic programming and planning. The inherent competitive aspect also can be used as a great motivator for students to design and implement their own AI systems.*

Keywords: *AI for games, Steer behavior, Path finding in 3d games and simulation*

I. INTRODUCTION

General game playing is the attempt to create a new generation of AI systems that are able to understand the rules of arbitrary games and then learn to play these games without human intervention. Unlike specialized systems like the chess program Deep Blue, a general game player cannot rely on algorithms designed in advance for specific games. Such a system rather requires a form of general intelligence that enables it to autonomously adapt to new and possibly radically different environments. General game-playing systems are a quintessential example of a new generation of software that end users can customize for their own specific tasks. This makes general game playing an interesting and challenging problem for AI, involving many fundamental issues such as reasoning, learning, planning and decision making. Consequently, general game playing can, and in fact should, be of interest to researchers in a variety of AI disciplines beyond conventional computer game playing. At the same time and for the same reasons, general game playing provides a new anchor for AI education as a unique and attractive framework for teaching multiple basic AI topics, such as problem solving by search, propositional and first-order logic, logic programming and planning. This paper attempts to give an overview of what has been achieved, and what lies ahead, some five years after the introduction of general game playing to the broad AI community through the inauguration of the annual AAAI General Game Playing Competition in 2005 [7]. Initiated and organized by Michael and members of his Stanford Logic Group (and endowed with the attractive purse of \$10,000), the competition quickly caught the interest of a number of researchers, the author included, from different backgrounds.

Yet the idea itself, that is, to build a system that can learn to play a whole array of games, has been around for over 40 years, going back to the French AI pioneer Jacques Pitarat, who wrote the first ever computer program that, in principle, could learn to play arbitrary chess-like board games by being given their rules [20]. Later general game-playing programs include [19, 12], but it required the aforementioned AAAI competition to spark broad interest in this problem as an AI Grand Challenge. Some five years later, an impressive number of research groups on general game playing have been established world-wide, including the German universities in Berlin, Bremen and Potsdam. Most of these groups develop their own player, but also there is an increasing number of researchers who are interested in specific aspects of general game playing, which does not require them to build a full-edged, competitive game-playing system. They all contribute to a fast growing body of literature on general game playing. Further indications that general game playing is a maturing research held are the establishment of a series of biennial IJCAI workshops on this topic in 2009;¹ the inclusion of "general game playing" among the standard keywords at the AAAI Conference; and the first ever special issue of a journal on this topic.² The first part of this paper will survey the research landscape of general game playing. We will see that a surprisingly broad range of classic AI fields have a role to play in general game playing. In each case we will show how existing approaches, methods and systems are contributing to the foundations for general game playing, to the improvement of the quality of existing general game-playing systems, and to the development of new methodologies. But even if your interest as an AI researcher does not lie in general game-playing systems themselves, they can be used as a non-trivial application for a broad range of more theoretically motivated AI research. We will also report on current research trends, identify some of the most pressing open questions, and look at the possibilities to gradually broadening today's concept of general game playing to involve even more aspects of AI.

In the second part of this paper, we will give an overview of how general game playing has entered AI education, either in form of an advanced AI graduate course, with lectures and tutorials but where the special focus lies on practical work; or as part of a general introductory course to AI. Examples for the former can be found in the curricula at Stanford University, Reykjavik University, Bremen University and TU Dresden, where it was held by the author for four consecutive years starting in winter 2006/7. Examples for the latter include the general introduction to AI for undergraduate students at the University of New South Wales to which the author contributed in spring 2011. We will show why general game playing provides an excellent angle for teaching a variety of basic

1. For GIGA'09 held in Pasadena see www.ru.is/faculty/yngvi/GIGA09; for GIGA'11 held in Barcelona see www.aiide.org/giga11.
2. KI/Künstliche Intelligenz, volume 25, Springer Verlag 2011. AI methods that also is a great motivator for students to design and implement their own AI systems.

Our overview will include a survey of freely available teaching aides including slides, tutorial questions and programming tools, for the benefit of potential instructors.

II. THE RESEARCH LANDSCAPE OF GENERAL GAME PLAYING

An outstanding characteristic of general game playing as an AI research topic is to involve a broad range of sub-disciplines with a focus on symbolic AI (as opposed to, say, RoboCup or the DARPA Grand Challenge for driverless vehicles). For this reason, general game playing has all the potential to become a rich source for interesting research problems in many different areas. As we will survey the research landscape, we will encounter several cases in which general game playing has been successfully used as an attractive [and challenging] application to demonstrate the viability of existing theories, methods, and systems. We will also see examples where the concept of general game playing has generated new research problems and solutions. Most importantly, it will become clear that it is not at all necessary to actually build a full-edged, competitive player in order to make an original and significant contribution to this research challenge. Yet another characteristic of general game playing research is to often concern the combination and integration of two or more theories and methods, which naturally leads to collaborations involving different AI sub-disciplines. In the relatively short time span since the first AAAI competition in 2005, at least four traditional AI disciplines have proved to be core aspects of research in general game playing:

1. Knowledge Representation
2. Search
3. Planning
4. Learning

In the remainder of this section, we will discuss in turn the role that each of these areas plays for general game playing: what interesting research problems they give rise to, which methods have been successfully applied, and what challenges lie ahead.

2.1 Knowledge Representation and Reasoning

General game playing requires a formal, symbolic language in which the rules of arbitrary games can be described to a system. The general Game Description Language (GDL) has been developed for that purpose [7]. It can be seen as a specific high-level action description language in the tradition of AI Planning languages like STRIPS or PDDL: game worlds are described with the help of individual features (e.g., the position of a piece on a game board), and moves are specified by their preconditions and effects. For example, the standard laws for chess include rules saying that a player whose turn it is can castle with the kingside rook under certain conditions, and as a result if white performs this action his king will move to square (g,1). In actual GDL, this looks something like the following, where GDL-specific keywords are printed in *italic* and variables are indicated by a leading *?*.

```
1 (<= ( legal ? player kingside_castle )
2   ( true ( turn ? player ) )
3   ( cancastletofkingside ? player ) )
4 (<= ( next ( cell ?g ?1 ? white_king ) )
5   ( does white kingside_castle ) )
```

This would be accompanied by an axiomatisation of the domain-specific predicate *can castle to kingside(x)* according to the standard rules for casting.

Game rules like these can be considered as a formal representation of knowledge for general game-playing systems. They are the source of a variety of interesting and challenging reasoning problems. Reasoning about actions. The use of GDL as the input language for general

game-playing systems immediately raises research questions such as: How to transform GDL descriptions in order to improve the efficiency of logical reasoning? How to apply existing techniques for reasoning about actions to GDL? How to extend or modify GDL for other applications of general intelligence? These problems are being extensively studied, and recent results include mappings of GDL into Propositional Automata [27], C++ code [34], and Binary Decision Diagrams [11]. GDL has been successfully embedded into a structure-rewriting formalism [9], the functional programming language OCAML [21], and the Situation Calculus [25]. Extensions of basic GDL have been developed to capture imperfect-information games [32] and coalition games [29]. However, a number of challenging problems remain, including

1. The development of a transformation of GDL into an efficient encoding that scales to large games for which full-edged grounding (i.e., propositionalisation) of the rules is impossible in practice;
2. Generalized reasoning methods for randomized, imperfect-information games;
3. The design of general problem description languages for other domains to provide the foundation for general-purpose agents, general-purpose robots, or for any other modern equivalent of the classic General Problem Solver in other AI disciplines.

Automated theorem proving. A key to successful general game playing is the ability of a system to use the right strategies and heuristics when playing a new game. In many cases this requires a system to analyze the game rules with the goal to establish game-specific properties that are not explicitly given but follow implicitly. Since games are specified in a logical language, this amounts to a theorem proving task. Successful approaches include methods for extracting general properties like symmetries [22], factorability [4, 26], or state invariants [33], and methods for identifying prevailing concepts such as mobility, boards, and pieces [13, 3]. Answer Set Solvers have been successfully deployed as general purpose theorem provers in this context [24, 31, 18]. Among the open problems are:

1. the identification of other common properties that a general game-playing system should determine in order to improve game-play;
2. The development and application of new non-propositional theorem proving techniques for games whose GDL rules cannot be fully grounded in practice.

Other KRR techniques. The logic-based formalization of game rules as the basis for general game playing implies that there are many as yet unexplored possibilities for other areas in Knowledge Representation to contribute to the general game playing effort. Some of a potentially large number of examples are 1. Reasoning About Uncertainty applied to games with randomised moves and where players have imperfect information; 2. Belief Revision employed, for example, to represent and revise beliefs about the opponents, as a way to build and maintain opponent models; 3. Spatial Reasoning to identify and reason about spatial relations in games.

2.2 Search

Once a general game-playing system is capable of computing legal moves and position updates from the game rules, it can search through the space of possible ways in which the game can proceed. Of course only very simple games allow a complete brute-force search in practice. Hence, one of the fundamental challenges for general game playing is to develop intelligent search techniques without knowing the particulars of the game to which they will be applied. The early general game-playing systems applied standard iterative deepening search in combination with automatically constructed heuristic evaluation function [3, 8, 13, 19, 23]. More recently, the use of Monte Carlo simulations has become main-stream in general game playing as it provides a way to achieve good performance without the need to generate any game-specific knowledge [2, 15, 18, 28]. The state-of-the-art search methods in existing general game-playing systems give rise to a number of interesting challenge problems. Monte Carlo Tree Search. Any effective use of random play outs in general game playing requires effective search-control methods. The standard algorithm is UCT (for: upper confidence bound applied to trees) [2], originally introduced to computer game playing by Kocsis and Szepesvári in 2006. A key to improving the performance of simulation-based general game playing is the ability to learn game-specific knowledge to effectively guide the random UCT simulations. Some such techniques have been described in [6, 10]. Among the open and challenging problems are the following.

1. The use of high-level game concepts such as those generated and used in heuristic evaluation function-based players for search control in Monte Carlo simulations;
2. The automation of the decision about the appropriate search technique and enhancement at runtime;
3. The generalization of simulation-based methods for general game playing to imperfect-information games, where a player may be highly uncertain about the current game state. Informed search. Informed search uses problem-specific knowledge to overcome the inefficiency of blind search. General game playing raises the challenge that any such knowledge needs to be automatically extracted. Existing methods for computing

symmetries [22] and factors [4, 26] can be used to cut down the search space. A recent method for computing general distance measures from the mere rules of a game can be used to automatically generate admissible search heuristics [16], but interesting problems remain, including

1. the use of other types of automatically extracted game-specific knowledge for effective search control;
2. the application of further informed search methods to general game playing, for example, local search algorithms.

2.3 Planning

Planning is closely related to general game playing, as both are instances of general problem solving, where the specifics of a problem are unknown until runtime. The game description language GDL can in fact be seen a special-purpose action description language that follows the tradition of planning languages. The latter, however, describe a problem from the perspective of a single agent, even in case of adversarial planning, and GDL generalizes this to the presence of other agents that have their own actions and goals. Reasoning about the intentions of the other players is the basis for Opponent Modeling alone of the crucial aspects in which general game playing goes beyond planning. A further difference is that planning is mainly concerned with solving a problem upfront, whereas much of general game playing is concerned with finding a good course of actions without being able to see all the way to the end of a game. Still, many aspects of planning are very relevant for general game playing, and already some results have been successfully transferred: In [5] it is shown how symbolic planning technology can be applied to explore the search space for both single-player and two-player games. Planning as Answer Set Programming has been successfully adapted to general game playing in [31]. The abovementioned automatic identification of symmetries [22] and decompositions [4, 26] and their use for game tree search are generalizations of similar results for planning, too. Some interesting open aspects at the boundary of planning and general game playing are the following.

1. the development of an automatic translations of GDL into an existing plan-ning language such as PDDL, which is challenging because GDL does not describe negative effects of actions explicitly (and instead appeals to the principle of negation-by-failure [14]);
2. the adaptation of successful planning heuristics, such as relaxation, to general game playing, which is challenging for similar reasons;
3. the generalisation of other planning techniques, e.g. hierarchical planning; 4. the converse application of techniques developed in general game playing, such as opponent modelling, for adversarial planning.

2.4 Learning

The very idea of general game playing is to build systems that automatically learn to master arbitrary new games. It is therefore justified to consider learning to be the core research area for this AI Challenge. Thus it may be somehow surprising that general AI learning methods are not (yet) dominating the state of the art in general game playing. The main reason seems to be the sheer difficulty of the problem for general-purpose learning algorithms, which is why existing, successful applications of learning techniques in general game playing focus on improving individual aspects of specific approaches: In [1], transfer learning has been applied to generalize knowledge from one game for other games. Neural networks have been used in [17] to improve automatically generated evaluation functions. In [6], statistical and reinforcement learning has been employed to provide search guidance in Monte Carlo simulations by biasing the random selection of moves. A variant of temporal difference learning has been developed in [10] to automatically construct features that are correlated with the outcome of a match. Also decision trees have been applied to evaluate states [30]. Despite these first attempts to employ learning techniques to improve the quality of various approaches, the application of learning in general game playing is still in its infancy. An obstacle may be that the current format of competitions is not favorable for many general learning methods, for two reasons: a comparably short amount of time to learn a new game (typically in the order of 10 minutes) and the fact that games are rarely repeated, if at all. General game playing as a principle, however, is more general than suggested by the special format of contemporary competitions. It can certainly be expected that learning will play an increasingly important role when systems have both more time| ideally using life-long learning and more data to learn from besides random play outs, including games they lost and games played by masters.

2.5 Further Relevant Research Areas

In addition to the four AI disciplines that today build the core of general game playing research, other areas are likely to eventually become relevant as the scope of general game playing is broadened.

Decision making. Current state-of-the-art research in general game playing focuses on games that are deterministic and in which players have complete knowledge of the state. As we progress towards the general case of randomized, im-perfect-information games, aspects of general decision making will become increasingly important, including utility functions, Markov decision processes, and game-theoretic solutions.

Natural language processing. Most end users of general game-playing systems will expect a better way than using GDL to input the rules of their own invented games. This could be achieved with the help of some controlled form of natural language for describing games. In most cases, this will include a process of disambiguation through a dialogue between the player and its user. This generalization seems particularly attractive because it could be met in a mostly modular fashion, where an existing general game player is coupled with a preprocessing natural language interface.

Game-playing robots. General game-playing software provides a great opportunity to make a relatively simple robotic system act smart. A robot arm capable of moving pieces on a board, coupled with a state-of-the-art player, can in principle learn to play arbitrary games with these pieces. Beyond this, interesting challenges are posed by game-playing robots that learn to recognize and manipulate new pieces, or mobile robots that can solve a whole array of new tasks formulated as games.

III. GENERAL GAME PLAYING IN AI EDUCATION

The fact that it addresses so many different aspects of AI makes general game playing a valuable tool for AI education, too, as it allows to teach a variety of methods on a single subject. A further advantage of general game playing is to naturally lead to practical assignments, where students can experiment with various techniques to address an interesting and challenging problem. Plus, there is the competitive aspect when students' players are pitted against each other for evaluation, which works as a great motivator. With today's available software tools and online resources for teaching general game playing, this is also much easier to organize than, say, a full-edged robotics laboratory. In this second part of the paper we give an overview of how general game playing can help with teaching the following core topics of AI:

1. Problem Solving (search, constraint satisfaction, game playing)
2. Knowledge and Reasoning (logic, inference, knowledge representation)
3. Planning

This will be complemented by an overview of teaching resources that are available for download, including slides, example tutorial assignments and software support for teaching general game playing.

3.1 Problem Solving

Problems solving by searching a solution space is often the first topic in an introductory AI course. Since general game playing is, in a certain sense, general problem solving, all of the following standard methods can be very nicely motivated and illustrated in this setting.

1. Problem formulation, that is, modeling a problem as a single-player game defining a state search space;
2. Uninformed search strategies like breadth-first, depth-first, iterative deepening etc., and methods for avoiding repeated states;
3. Informed search using heuristic functions;
4. Constraint satisfaction as a special kind of general problem formulation;
5. Adversarial search using the minimax algorithm and alpha-beta pruning.

In fact, many of the standard examples used in textbooks to illustrate these methods are classic games, such as the 15-puzzle for state-based search, the 8-queens problem for constraint solving, and tic-tac-toe for minimax and alpha beta pruning.

3.2 Knowledge and Reasoning

General game playing requires to axiomatise the rules of games, which makes it an ideal case study for learning how to use symbolic logic to represent knowledge. Through the problem of describing a game formally, students learn to choose appropriate symbols for representing the various elements of a game (such as state features and moves) and to formulate in (propositional or first-order) logic laws that are given in natural language. A basic general game-playing system

must be able to reason about games described in GDL in order to determine its legal moves and their effects etc., which can be used to motivate studying various inference methods. In this way, general game playing allows to address all of the following topics:

1. Propositional logic, using simple games with a small state space;
2. Resolution to draw inferences from game rules;
3. First-order logic for games with rules that are naturally described using variables and quantification;
4. Logic programming as the foundation for a basic general game player.

3.3 Planning

We have already argued earlier, in Section 2.3, why general game playing is closely related to planning and in fact can be seen as a generalization thereof. Hence, all aspects of planning can naturally be taught under the umbrella of general game playing, too:

1. Classic Planning for solving single-player games;
2. Continuous Planning for playing single-player games;
3. Conditional Planning for nondeterministic and partially observable games;
4. Multiagent- and Adversarial Planning for cooperative/competitive games.

3.4 Teaching Resources

The author wants to stress that none of the above is meant to suggest conveying to students the impression that general game playing is the only reason for studying problem solving, knowledge representation or planning. Rather, the idea is to provide a domain to which students can repeatedly return in order to practice a new theory and experiment with a new method. Thus there are two principled ways to use general game playing in AI education, both of which the author and a number of other instructors at various universities have successfully used in the past:

1. A specialized graduate course on general game playing, the ultimate goal of which is for students to develop their own players and where, along the way, they learn to apply the various relevant AI methods. Experiences with such a course at Stanford University and TU Dresden over several years suggest that the right size for such a course is about 15{30 participants with groups of 2-4 (ideally, 3) students designing one player.
2. Integrating general game playing into a general introductory AI course, both on the undergraduate and the graduate level.

For either type of course, the website www.general-game-playing.de is a rich source for prospective instructors of free supporting material:

1. Complete sets of slides from several universities worldwide for full-edged general game playing courses and for parts of introductory AI courses.
2. Sample tutorial questions and assignments.
3. Basic players in several programming languages available for download that students can be provided with as a starting point.
4. Software to run a student competition in class.

Any general introductory course to AI can in fact profit from the available tools even without reference to general game playing: Instructors can use the general game control software to let students experiment with any concrete game or agent domain, like e.g. the famous Wumpus World, for which formalization in GDL exists.

The author welcomes any additional material or pointers to new courses on general game playing.

REFERENCES

- [1] Barwood, H. & Falstein, N. 2002. .More of the 400: Discovering Design Rules.. Lecture at Game Developers Conference, 2002. Available online at: <http://www.gdconf.com/archives/2002/halfbarwood.ppt>
- [2] Church, D. 1999. .Formal Abstract Design Tools.. Game Developer, August 1999. San Francisco, CA: CMP Media. Available online at: <http://www.gamasutra.com/features/19990716/designftoolsf01.htm>
- [3] Hunicke, R. 2004. .AI Babysitter Elective.. Lecture at Game Developers Conference Game Tuning Workshop, 2004. In LeBlanc et al., 2004a. Available online at: <http://algorithmancy.8kindsoffun.com/GDC2004/AITutorial5.ppt> LeBlanc, M., ed. 2004a. .Game Design and Tuning
- [4] Workshop Materials., Game Developers Conference 2004. Available online at: <http://algorithmancy.8kindsoffun.com/GDC2004/>
- [5] LeBlanc, M. 2004b. .Mechanics, Dynamics, Aesthetics: A Formal Approach to Game Design.. Lecture at Northwestern University, April 2004.

About Authors

Pranav Kumar Pathak

I have done Masters from Sardar Patel University, India. Currently pursuing Ph.D in Computer Science (AI) from Sardar Patel University, India. I have also done MBA IT from Manipal University, India. Currently I am working as Lecturer in Computer Science Department at College of Computer and Information Sciences in King Saud University, Kingdom of Saudi Arabia. I have presented several papers at international Journals and conferences.

Dr. Dipti M.Shah

I have done Ph.D in Computer Science from Sardar Patel University, India. Currently I am working as Associate Professor in Computer Science Department at G.H. Patel Pg Department Of Computer Science & Technology V.V.Nagar. I have presented several papers at international Journals and conferences.